

7(a)	One from: - A* algorithm - Dijkstra's algorithm
7(b)	One mark for each mark point (Max 5) MP1 Deep learning is a branch of machine learning MP2 It makes use of artificial neural networks MP3 ... to work in a similar manner to the human/biological brain // ... neural networks can be trained to process data MP4 It is a structure consisting of many interconnected layers MP5 ... input, hidden and output // the input layer feeds the first hidden layer, the last hidden layer feeds the output MP6 ... gradually breaking down and analysing the data to find (hidden) patterns MP7 Deep learning's feature extraction/learning involves how to detect features/categories from data MP8 ... and can involve different types of supervision MP9 It works well with large data sets MP10 Deep learning uses back propagation of errors to adjust parameters/weightings to refine its results / improve its accuracy.

9(a)(i)	One from: e.g. - Virtual assistants // AI Agents - Chatbots - Composing music - Self-driving cars - Natural language processing - Fraud detection - Image colouring // photograph enhancement - Visual/Image recognition/facial recognition - Text mining - Computer assisted translation - Large language models (LLMs) - Playing chess - Grading student exams - Weather forecasting
9(a)(ii)	Use / introduce more hidden layers
9(b)	One mark per mark point (Max 4) MP1 Initial outputs are compared to expected outputs MP2 ... weightings are adjusted to minimise the difference between actual and expected outputs MP3 Calculus is used to find the error gradient in the obtained outputs MP4 ... the results are fed back into the neural network MP5 ... weightings of each neuron / node are adjusted as a result of the feedback MP6 ... the process repeats until results are more accurate

3(a)	1 mark each to max 4 e.g. - Uses image recognition // optical character recognition - to analyse a photograph/pixels to identify the location of characters within an image - in order to convert the pattern of pixels into individual characters - Uses natural language - to combine the characters into words - and compare the words to database to identify a match in the language of the writing - it may also predict what is missing / incomplete / from what has been read - then compare the words to a database of the user's chosen language to find a match and translate - Uses text-to-speech - to generate the audio waveforms corresponding to the words
3(b)	1 mark each to max 2 e.g. - It can help people with visual impairment to identify which products to purchase - It can help people who cannot understand the language - It can help to overcome learning / reading difficulties
3(c)(i)	1 mark each to max 3 - The program is usually sold for a fee - Users have no access to source code - It protects the intellectual rights of the programmer / developer - Users cannot legally alter or share the program - There are limitations on the use of the software e.g. number of times it can be installed - Support and maintenance is usually provided e.g. updates, bug fixes etc.
3(c)(ii)	1 mark each to max 3 e.g. - The source code would be accessible therefore it could be changed - which might lead to the program outputting an incorrect response - so people could purchase the wrong item - which might have serious consequences // by example e.g. allergies

6(a)	One mark per point (Max 2) MP1 A graph uses vertices/nodes to identify/represent entities such as destinations, people, etc MP2 Edges are used to connect nodes and can represent possible paths between them // a path is the list of nodes connected by edges between two given nodes MP3 Nodes/edges can be labelled/weighted, and this is a weighting that can be applied and used in the context of the application MP4 A cycle is a list of nodes that return to the same node.
6(b)	One mark per mark point (Max 4) MP1 Supervised learning uses labelled data // Unsupervised learning makes use of unlabelled data. MP2 Labelled data means that known outcomes are applied to specific inputs to help the AI predict outcomes. MP3 Supervised learning requires initial human input/training // Unsupervised learning does not require human input/training. MP4 With unlabelled data in unsupervised learning, outcomes are not known MP5 ... the AI has to search for hidden patterns/structures/clusters MP6 ... within the data in order to predict outcomes.

10(a)	To find the path between two points on a graph using the algorithm.
10(b)	One mark per point (Max 2) MP1 A* tries to find a better path (between two points) by using a heuristic function // A* finds the adjacent route with the shortest path and continues this until the destination is reached MP2 ... Dijkstra's just explores all possible routes. MP3 The heuristic function on the A* algorithm gives priority to nodes that are supposed to be better than others / less costly than others. MP4 Dijkstra's algorithm cannot work with negative values/weights //A* algorithm can work with negative values/weights.
10(c)	One mark for each mark point (Max 3) MP1 Unsupervised learning uses algorithms to analyse / cluster MP2 ... unlabelled data sets MP3 They discover hidden patterns / data groupings / clusters without the need for human intervention. MP4 It is able to discover similarities and differences in data / information.

6	1 mark for each bullet point (max 4) e.g. • It uses image recognition • The pixels of each image from the camera are stored • ... and matched to the expected shape/size/colour of registration number • It uses optical character recognition • The pixels within the registration number area identified are analysed • ... and compared to expected characters/letters/numbers
---	---